arXiv:1607.05112v1 [cs.DS] 18 Jul 2016

Minimum cycle and homology bases of surface embedded graphs*

Glencora Borradaile† Erin Wolf Chambers‡ Kyle Fox§ Amir Nayyeri†

July 19, 2016

Abstract

We study the problems of finding a minimum cycle basis (a minimum weight set of cycles that form a basis for the cycle space) and a minimum homology basis (a minimum weight set of cycles that generates the 1-dimensional ($\mathbb{Z}_2$)-homology classes) of an undirected graph embedded on a surface. The problems are closely related, because the minimum cycle basis of a graph contains its minimum homology basis, and the minimum homology basis of the 1-skeleton of any graph is exactly its minimum cycle basis.

For the minimum cycle basis problem, we give a deterministic $O(n^\omega + 2^{2g}n^2 + m)$-time algorithm for graphs embedded on an orientable surface of genus g. The best known existing algorithms for surface embedded graphs are those for general graphs: an $O(m^\omega)$ time Monte Carlo algorithm [2] and a deterministic $O(nm^2/\log n + n^2m)$ time algorithm [31]. For the minimum homology basis problem, we give a deterministic $O((g+b)^3 n \log n + m)$-time algorithm for graphs embedded on an orientable or non-orientable surface of genus g with b boundary components, assuming shortest paths are unique, improving on existing algorithms for many values of g and n. The assumption of unique shortest paths can be avoided with high probability using randomization or deterministically by increasing the running time of the homology basis algorithm by a factor of $O(\log n)$.

*A preliminary version of this work was presented at the 32nd Annual International Symposium on Computational Geometry [4]. This material is based upon work supported by the National Science Foundation under grants CCF-12-52833, CCF-10-54779, IIS-13-19573, CCF-11-61359, IIS-14-08846, CCF-15-13816, CCF-1617951, and IIS-14-47554; by an ARO grant W911NF-15-1-0408; and by Grant 2012/229 from the U.S.-Israel Binational Science Foundation.

†Oregon State University, Corvallis, OR; {glencora,nayyeria}@eecs.orst.edu.

‡St. Louis University, St. Louis, MO; echambe5@slu.edu.

§Duke University, Durham, NC; kylefox@cs.duke.edu.

1 Introduction

1.1 Minimum cycle basis

Let $G = (V, E)$ be a connected undirected graph with n vertices and m edges. We define a ***cycle*** of G to be a subset $E' \subseteq E$ where each vertex $v \in V$ is incident to an even number of edges in E'. The ***cycle space*** of G is the vector space over cycles in G where addition is defined as the symmetric difference of cycles' edge sets. It is well known that the cycle space of G is isomorphic to $\mathbb{Z}_2^{m-n+1}$; in particular, the cycle space can be generated by the fundamental cycles of any spanning tree of G. A ***cycle basis*** is a maximal set of independent cycles. A ***minimum cycle basis*** is a cycle basis with a minimum number of edges (counted with multiplicity) or minimum total weight if edges are weighted[1]. Minimum cycle bases have applications in many areas such as electrical circuit theory [10, 28], structural engineering [9], surface reconstruction [35], and the analysis of algorithms [30].

Sets of independent cycles form a matroid, so the minimum cycle basis can be computed using the standard greedy algorithm. However, there may be an exponential number of cycles in G. Horton [25] gave the first polynomial time algorithm for the problem by observing that every cycle in the minimum cycle basis is the fundamental cycle of a shortest path tree. Several other algorithms have been proposed to compute minimum cycle bases in general graphs [2,3,12,20,27,31]. The fastest of these algorithms are an $O(m^\omega)$ time Monte Carlo randomized algorithm of Amaldi *et al.* [2] and an $O(nm^2/\log n + n^2 m)$ time deterministic algorithm of Mehlhorn and Michail [31]. Here, $O(m^\omega)$ is the time it takes to multiply two $m \times m$ matrices using fast matrix multiplication.

For the special case of *planar graphs*, faster algorithms are known. Hartvigsen and Mardon [23] observed that the cycles in the minimum cycle basis nest, implicitly forming a tree; in fact, the edges of each cycle span an s, t-minimum cut between two vertices in the dual graph, and the Gomory-Hu tree [21] of the dual graph is precisely the tree of minimum cycle basis in the primal. Hartvigsen and Mardon [23] gave an $O(n^2 \log n)$ time algorithm for the minimum cycle basis problem in planar graphs, and Amaldi *et al.* [2] improved their running time to $O(n^2)$. Borradaile, Sankowski, and Wulff-Nilsen [6] showed how to compute an *oracle* for the minimum cycle basis and dual minimum cuts in $O(n \log^4 n)$ time that is able to report individual cycles or cuts in time proportional to their size. Borradaile *et al.* [5] recently generalized the minimum cut oracle to graphs embeddable on an orientable surface of genus g. Their oracle takes $2^{O(g^2)} n \log^3 n$ time to construct (improving upon the original planar oracle by a factor of $\log n$). Unfortunately, their oracle does not help in finding the minimum cycle basis in higher genus graphs, because there is no longer a bijection between cuts in the dual graph and cycles in the primal graph.

That said, it is not surprising that the cycle basis oracle has not been generalized beyond the plane. While cuts in the dual continue to nest in higher genus surfaces, cycles do not. In fact, the minimum cycle basis of a toroidal graph must *always* contain at least one pair of crossing cycles, because any cycle basis must contain cycles which are topologically distinct. These cycles must represent different *homology classes* of the surface.

1.2 Minimum homology basis

Given a graph G embedded in a surface Σ of genus g with b boundary components, the homology of G is an algebraic description of the topology of Σ and of G's embedding. In this paper, we focus on one-dimensional cellular homology over the finite field $\mathbb{Z}_2$. Homology of this type allows for

[1] There is a notion of minimum cycle bases in directed graphs as well, but we focus on the undirected case in this paper.

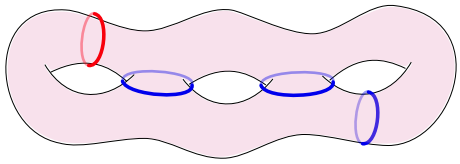

Figure 1. Two homologous cycles, one shown in red and the other in blue.

simplified definitions. We say a cycle η is ***null-homologous*** if η is the boundary of a subset of faces. Two cycles η and η' are ***homologous*** or in the same ***homology class*** if their symmetric difference $\eta \oplus \eta'$ is null-homologous. Let $\beta = g - \max\{b-1, 0\}$ if Σ is *non-orientable* (contains a subset homeomorphic to the Möbius band), and let $\beta = 2g - \max\{b-1, 0\}$. The homology classes form a vector space isomorphic to $\mathbb{Z}_2^\beta$ known as the ***homology space***. A ***homology basis*** of G is a set of β cycles in linearly independent homology classes, and the ***minimum homology basis*** of G is the homology basis with either the minimum number of edges or with minimum total weight if edges of G are weighted.

Erickson and Whittlesey [18] described an $O(n^2 \log n + gn^2 + g^3 n)$ time algorithm for computing the minimum homology basis for orientable Σ without boundary. Like Horton [25], they apply the greedy matroid basis algorithm to a set of $O(n^2)$ candidate cycles. Alternatively, a set of 2^β candidate cycles containing the minimum homology basis can be computed easily by applying the algorithms of Italiano *et al.* [26] or Erickson and Nayyeri [17] for computing the minimum homologous cycle in any specified homology class. These algorithms take $g^{O(g)} n \log\log n$ and $2^{O(g)} n \log n$ time respectively. While Erickson and Whittlesey [18] do not explicitly consider the case, all three results mentioned above can be extended to surfaces with boundary. Similarly, the algorithms of Erickson and Whittlesey [18] and Erickson and Nayyeri [17] can be applied to compute the minimum homology basis for non-orientable Σ, even though they only consider orientable surfaces explicitly. Dey, Sun, and Wang [13] generalized the results above to arbitrary simplicial complexes, and Busaryev *et al.* [7] improved the running time of their generalization from $O(n^4)$ to $O(n^\omega + n^2 g^{\omega-1})$. Note that all of the algorithms above either take quadratic time in n (or worse) or they have exponential dependency on g. In contrast, it is well understood how to find exactly one cycle of the minimum homology basis of G in only $O(g^2 n \log n)$ time assuming orientable Σ, because the minimum weight non-separating cycle will always be in the basis [8, 16].

1.3 Our results

We describe new algorithms for computing the minimum cycle basis and minimum homology basis of the graph G. Our algorithm for minimum cycle basis requires G be embedded on an orientable Σ, but it is deterministic and runs in $O(n^\omega + 2^{2g} n^2 + m)$ time, matching the running time of the randomized algorithm of Amaldi *et al.* [2] when g is sufficiently small. Our algorithm for minimum homology basis works in orientable or non-orientable Σ, is also deterministic, and it runs in $O((g+b)^3 n \log n + m)$ time assuming shortest paths are unique. The assumption of unique shortest paths is only necessary to use the multiple-source shortest path data structure of Cabello, Chambers, and Erickson [8]. It can be avoided with high probability by using randomization or deterministically by increasing the running time of our algorithm by a factor of $O(\log n)$ [8]. For simplicity, we will assume shortest paths are unique during the presentation of our minimum homology basis algorithm. In any case, ours is the first algorithm for minimum homology basis that has a running time simultaneously near-linear in n and polynomial in g.

At a high level, both of our algorithms are based on the $O(nm^2 + n^2 m \log n)$ time algorithm of Kavitha *et al.* [27] who in turn use an idea of de Pina [12]. We compute our basis cycles one by one.

Over the course of the algorithm, we maintain a set of *support vectors* that form the basis of the subspace that is *orthogonal* to the set of cycles we have already computed. Every time we compute a new cycle, we find one of minimum weight that is *not* orthogonal to a chosen support vector S, and then update the remaining support vectors so they remain orthogonal to our now larger set of cycles. Using the divide-and-conquer approach of Kavitha *et al.* [27], we are able to maintain these support vectors in only $O(n^\omega)$ time total in our minimum cycle basis algorithm and $O(g^\omega)$ time total in our minimum homology basis algorithm. Our approaches for picking the minimum weight cycle not orthogonal to S form the more technically interesting parts of our algorithms and are unique to this work.

For our minimum cycle basis algorithm, we compute a collection of $O(2^{2g}n)$ cycles that contain the minimum cycle basis and then partition these cycles according to their homology classes. The cycles within a single homology class nest in a similar fashion to the minimum cycle basis cycles of a planar graph. Every time we compute a new cycle for our minimum cycle basis, we walk up the 2^{2g} trees of nested cycles and find the minimum weight cycle not orthogonal to S in $O(n)$ time per tree. Overall, we spend $O(2^{2g}n^2)$ time finding these cycles; if any improvement is made on the time it takes to update the support vectors, then the running time of our algorithm as a whole will improve as well.

Our minimum homology basis algorithm uses a covering space called the cyclic double cover. As shown by Erickson [16], the cyclic double cover provides a convenient way to find a minimum weight closed walk γ crossing an arbitrary non-separating cycle λ an odd number of times. We extend his construction so that we may consider not just one λ but any arbitrarily large collection of cycles. Every time we compute a new cycle in our minimum homology basis algorithm, we let S determine a set of cycles that must be crossed an odd number of times, build the cyclic double cover for that set, and then compute our homology basis cycle in $O((g+b)gn\log n)$ time by computing minimum weight paths in the covering space[2].

The rest of the paper is organized as follows. We provide more preliminary material on surface embedded graphs in Section 2. In Section 3, we describe a characterization of cycles and homology classes using binary vectors. These vectors are helpful in formally defining our support vectors. We give a high level overview of our minimum cycle basis algorithm in Section 4 and describe how to pick individual cycles in Section 5. Finally, we give our minimum homology basis algorithm in Section 6.

2 Preliminaries

We begin with an overview of graph embeddings on surfaces. For more background, we refer readers to books and surveys on topology [24, 33], computational topology [14, 37], and graphs on surfaces [11, 32].

A ***surface*** or 2-manifold with boundary Σ is a compact Hausdorff space in which every point lies in an open neighborhood homeomorphic to the Euclidean plane or the closed half plane. The ***boundary*** of the surface is the set of all points whose open neighborhoods are homeomorphic to the closed half plane. Every boundary component is homeomorphic to the circle. A ***cycle*** in the surface Σ is a continuous function $\gamma : S^1 \to \Sigma$, where S^1 is the unit circle. Cycle γ is called ***simple*** if γ is injective. A ***path*** in surface Σ is a continuous function $p : [0,1] \to \Sigma$; again, path p is simple

[2] In addition to the above results, we note that it is possible to improve the $g^{O(g)}n\log\log n$ time algorithm for minimum homology basis based on Italiano *et al.* [26] so that it runs in $2^{O(g)}n\log\log n$ time. However, this improvement is a trivial adaption of techniques used by Fox [19] to get a $2^{O(g)}n\log\log n$ time algorithm for minimum weight non-separating and non-contractible cycle in undirected graphs. We will not further discuss this improvement in our paper.

if p is injective. A ***loop*** is a path p such that $p(0) = p(1)$; equivalently, it is a cycle with a designated basepoint. The ***genus*** of the surface Σ, denoted by g, is the maximum number of disjoint simple cycles $\gamma_1, \ldots, \gamma_g$ in Σ such that $\Sigma \setminus (\gamma_1 \cup \cdots \cup \gamma_g)$ is connected. Surface Σ is ***non-orientable*** if it contains a subset homeomorphic to the Möbius band; otherwise, it is ***orientable***.

The ***embedding*** of graph $G = (V, E)$ is a drawing of G on Σ which maps vertices to distinct points on Σ and edges to internally disjoint simple paths whose endpoints lie on their incident vertices' points. A ***face*** of the embedding is a maximally connected subset of Σ that does not intersect the image of G. An embedding is ***cellular*** if every face is homeomorphic to an open disc; in particular, every boundary component must be covered by (the image of) a cycle in G. These boundary cycles must be vertex-disjoint. Embeddings can be described combinatorially using a rotational system and a signature. The ***rotation system*** describes the cyclic ordering of edges around each vertex. The ***orientation signature*** $\text{sig} : E \to \{0, 1\}$ is a function that assigns to each edge e a bit. Value $\text{sig}(e) = 0$ if the cyclic ordering of e's endpoints are in the same direction; otherwise, $\text{sig}(e) = 1$. Abusing notation, we denote the orientation signature of a cycle η (in G) as $\text{sig}(\eta)$ and define it as the exclusive-or of its edges' orientation signatures. If $\text{sig}(\eta) = 1$, we say η is ***one-sided***. Otherwise, we say η is ***two-sided***. Surface Σ is orientable if and only if every cycle of G is two-sided.

Let F be the set of faces in G. Let n, m, ℓ, and b be the number of vertices, edges, faces, and boundary components of G's embedding respectively. The ***Euler characteristic*** $\boldsymbol{\chi}$ of Σ is $2-2g-b$ if Σ is orientable and is $2-g-b$ otherwise. By Euler's formula, $\chi = n-m+\ell$. Embedded graphs can be ***dualized***: G^* is the graph embedded on the same surface, with a vertex in G^* for every face *and boundary component* in G and a face in G^* for every vertex of G. We refer to the dual vertices of boundary components as ***boundary dual vertices***. Two vertices in G^* are adjacent if the corresponding faces/boundary components are separated by an edge in G. We generally do not distinguish between edges in the primal and dual graphs. We assume Σ contains at least one boundary component as it does not affect the homology of Σ to remove a face when there are no boundary to begin with. In particular, this assumptions simplifies the definition of β as given in the introduction so $\beta = g - b + 1$ if Σ is non-orientable and $\beta = 2g - b + 1$ otherwise.

A ***spanning tree*** of the graph G is a subset of edges of G which form a tree containing every vertex. A ***spanning coforest*** is a subset of edges which form a forest in the dual graph with exactly b components, each containing one dual boundary vertex. A ***tree-coforest decomposition*** of G is a partition of G into 3 edge disjoint subsets, (T, L, C), where T is a spanning tree of G, C is a spanning coforest, and L is the set of leftover edges $E \setminus (T \cup C)$ [15, 17]. Euler's formula implies $|L| = \beta$.

A ***w, w'-path*** p (in G) is an ordered sequence of edges $\{u_1v_1, u_2v_2, \ldots, u_kv_k\}$ where $w = u_1$, $w' = v_k$, and $v_i = u_{i+1}$ for all positive $i < k$; a ***closed path*** is a path which starts and ends on the same vertex. A path is ***simple*** if it repeats no vertices (except possibly the first and last). We sometimes use ***simple cycle*** to mean a simple closed path. A path in the dual graph G^* is referred to as a ***co-path*** and a cycle in the dual is referred to as a ***co-cycle***. Simple paths and cycles in the dual are referred to as simple co-paths and co-cycles respectively. Every member of the minimum cycle basis (and subsequently the minimum homology basis) is a simple cycle [25]. We let $\sigma(u, v)$ denote an arbitrary shortest (minimum weight) u, v-path in G. Let $p[u, v]$ denote the ***subpath*** of p from u to v. Given a u, v-path p and a v, w-path p', let $p \cdot p'$ denote their concatenation. Two paths p and p' ***cross*** if their embeddings in Σ cannot be be made disjoint through infinitesimal perturbations; more formally, they cross if there is a maximal (possibly trivial) common subpath p'' of p and p' such that, upon contracting p'' to a vertex v, two edges each of p and p' alternate in their embedded around v. Two closed paths cross if they have subpaths which cross.

Let γ be a closed path in G that does not cross itself. We define the operation of ***cutting***

along γ and denote it $G \not\downarrow \gamma$. Graph $G \not\downarrow \gamma$ is obtained by cutting along γ in the drawing of G on the surface, creating two copies of γ. If $\text{sig}(\gamma) = 0$, then the two copies of γ each form boundary components in the cut open surface. Otherwise, the two copies of γ together form a single closed path that is the concatenation of γ to itself at both ends; the single closed path forms a single boundary component. Likewise, given a simple path σ in G, we obtain the graph $G \not\downarrow \sigma$ by cutting along σ, creating two interiorly disjoint copies connected at their endpoints. The cut open surface has one new boundary component bounded by the copies of σ.

Let F' be a collection of faces and boundary components. Let $\partial F'$ denote the boundary of F', the set of edges with exactly one incident face or boundary component in F'. We sometimes call F' a ***cut*** of G^* and say $\partial F'$ ***spans*** the cut. A co-path p with edge $uv \in \partial F'$ ***crosses*** the cut at uv.

Finally, let w and w' be two bit-vectors of the same length. We let $\langle w, w' \rangle$ denote the ***dot product*** of w and w', defined by the exclusive-or of the products of their corresponding bits.

2.1 Sparsifying G

We assume $g = O(n^{1-\varepsilon})$ for some constant $\varepsilon > 0$; otherwise, our algorithms offer no improvement over previously known results. Because boundary cycles are vertex-disjoint, we also have $b = O(n)$. Here, we describe a procedure to preprocess G so it contains a linear (in n) number of edges and faces as well. Namely, we modify G so it contains no faces of degree 2 or 1; Euler's formula immediately limits the number of edges and faces in the modified graph to $O(n)$. We begin by computing all pairs shortest paths. Observe, if we simply remove parallel edges from G, then we only reduce the minimum genus of any surface into which we can embed G. In particular, the graph would have $O(n)$ edges left after removing all parallel edges. All pairs shortest paths can be computed in $O(n^2 \log n + m)$ time.

Now, consider the graph G as given in the input. We iteratively perform the following procedure until every face has degree 3 or greater or our graph is one of a constant number of easy cases. In each iteration, we add at most one cycle to the minimum cycle basis or minimum homology basis. Let f be a face of degree 1 or 2. If f has degree 1, then it is bounded by a null-homologous loop e in G. We add $\{e\}$ to the minimum cycle basis, because it is the cheapest cycle containing e, but we do not add it to the minimum homology basis. If G consists entirely of e, we terminate; otherwise we remove e and f from the graph and continue with the next iteration. If f has degree 2, then it is either bounded by two distinct edges e and e' or bounded twice by a single edge. In the latter case, graph G must be the path of length 1 embedded in the sphere or it is a single vertex and non-null homologous loop embedded in the projective plane (the non-orientable surface of genus 1). If it is the path in the sphere, we add nothing to the minimum cycle basis and minimum homology basis, and we terminate. If it is a loop in the projective plane, we add it to both the minimum cycle basis and the minimum homology basis and terminate. Now suppose f is bounded by distinct faces e and e', and let e have less weight than e' without loss of generality. Edge e' belongs to cycle $\{e, e'\}$, so it belongs to some cycle of the minimum cycle basis. Let σ be the shortest path between the endpoints of e'. We add $\sigma \cdot e'$ to the minimum cycle basis. No other cycle in the minimum cycle basis contains e', because it would always be at least as cheap to include e in the cycle instead. Also, no cycle of the minimum homology basis contains e', because it would always be at least as cheap to include e and $\{e, e'\}$ itself is null-homologous. We remove e' and f from the graph and continue with the next iteration.

Each iteration is done in constant time, and there are at most m iterations of the above algorithm. Therefore, the preprocessing procedure takes $O(n^2 \log n + m)$ time total. We assume for the rest of that paper that $m = O(n)$ and $\ell = O(n)$.

3 Cycle and Homology Signatures

We begin the presentation of our algorithms by giving a characterization of cycles and homology classes using binary vectors. These vectors will be useful in helping us determine which cycles can be safely added to our minimum cycle and homology bases. Let (T, L, C) be an arbitrary tree, coforest decomposition of G; set L contains exactly β edges $e_1, \ldots, e_\beta$. For each index $i \in \{1, \ldots, \beta\}$, graph $C \cup \{e_i\}$ contains a unique simple co-cycle or a unique simple co-path between distinct dual boundary vertices. Let p_i denote this simple co-cycle or co-path. Let $f_{\beta+1}, \ldots, f_{m-n+1}$ denote the $m - n + 1 - \beta = \ell$ faces of G, and for each index $i \in \{\beta + 1, \ldots, m - n + 1\}$, let p_i denote the simple co-path from f_i to the dual boundary vertex in f_i's component of C.

For each edge e in G, we define its ***cycle signature*** $[e]$ as an $(m - n + 1)$-bit vector whose ith bit is equal to 1 if and only if e appears in p_i. The cycle signature $[\eta]$ of any cycle η is the bitwise exclusive-or of the signatures of its edges. Equivalently, the ith bit of $[\eta]$ is 1 if and only if η and p_i share an odd number of edges. Similarly, for each edge e in G, we define its ***homology signature*** $[e]_h$ as a β-bit vector whose ith bit is equal to 1 if and only if e appears in p_i. The homology signature of cycles is defined similarly.

The following lemma is immediate.

Lemma 3.1. *Let η and η' be two cycles. We have $[\eta \oplus \eta'] = [\eta] \oplus [\eta']$ and $[\eta \oplus \eta']_h = [\eta]_h \oplus [\eta']_h$.*

Let ζ_i denote the unique simple cycle in $T \cup \{e_i\}$. The following lemma helps us explain the properties of cycle and homology signatures.

Lemma 3.2. *The set of cycles $\{\zeta_1, \ldots, \zeta_\beta\}$ form a homology basis.*

Proof: We prove that the cycles lie in independent homology classes by showing the symmetric difference of any non-empty subset of $\{\zeta_1, \ldots, \zeta_\beta\}$ is not null-homologous. Suppose to the contrary that there exists a non-empty $\Upsilon \subseteq \{\zeta_1, \ldots, \zeta_\beta\}$ such that $\bigoplus_{\eta \in \Upsilon} \eta = \partial F'$ for some subset of faces $F' \subseteq F$, where $\bigoplus$ is the symmetric difference of its operands. Let $\zeta_i \in \Upsilon$ be an arbitrary member of the subset. Co-path p_i shares exactly one edge with ζ_i, and it shares no edges with any other $\eta \in \Upsilon$. In particular, p_i crosses dual cut F' an odd number of times. Therefore, p_i cannot be a co-cycle. Further, p_i cannot be a co-path between two distinct dual boundary vertices, because exactly one of those two vertices would have to lie inside F', a contradiction on F' only containing faces. We conclude Υ cannot exist and the cycles $\{\zeta_1, \ldots, \zeta_\beta\}$ do lie in independent homology classes. □

Let w be an arbitrary $(m-n+1)$-bit vector. We construct a cycle η_w to demonstrate how cycle and homology signatures provide a convenient way to distinguish between cycles and their homology classes. Let $\Upsilon \subseteq \{\zeta_1, \ldots, \zeta_\beta\}$ be the subset of basis cycles containing exactly the cycles ζ_i such that the ith bit of w is equal to 1. Similarly, let $F' \subseteq F$ be the subset of faces such that face $f_i \in F'$ if and only the ith bit of w is equal to 1. Let $\eta_w = \bigoplus_{\eta \in (\Upsilon \cup \{\partial F'\})} \eta$.

Lemma 3.3. *We have $[\eta_w] = w$.*

Proof: Let $i \in \{1, \ldots, m - n + 1\}$, and let p_i be the co-path as defined above. Suppose $i \in \{1, \ldots, \beta\}$. Co-path p_i crosses cut F' an even number of times. If bit i in w is set to 1, then p_i shares exactly one edge of $\bigoplus_{\eta \in \Upsilon} \eta$ by construction, and it must share an odd number of edges with η_w as well. If bit i in w is set to 0, then p_i shares no edges with $\bigoplus_{\eta \in \Upsilon} \eta$, and it must share an even number of edges with η_w.

Now, suppose $i \in \{\beta+1, \ldots, m-n+1\}$. Co-path p_i shares no edges with $\bigoplus_{\eta\in\Upsilon}\eta$. If i is set to 1, then $f_i \in F'$ and p_i crosses cut F' an odd number of times. Therefore, it shares an odd number of edges with η_w. If i is set to 0, then $f_i \notin F'$, and p_i crosses cut F' an even number of times, sharing an even number of edges with η_w. □

Corollary 3.4. *Let η and η' be two cycles. We have $\eta = \eta'$ if and only if $[\eta] = [\eta']$.*

Observe that the homology class of η_w is entirely determined by the first β bits of w. We immediately obtain an alternative (and more combinatorially inspired) proof of the following corollary of Erickson and Nayyeri [17].

Corollary 3.5 (Erickson and Nayyeri [17, Corollary 3.3]). *Two cycles η and η' are homologous if and only if $[\eta]_h = [\eta']_h$.*

Corollary 3.6. *Cycle signatures are an isomorphism between the cycle space and $\mathbb{Z}_2^{m-n+1}$, and homology signatures are an isomorphism between the first homology space and $\mathbb{Z}_2^{2g}$.*

4 Minimum Cycle Basis

We now describe our algorithm for computing a minimum cycle basis. We assume without loss of generality that surface Σ contains exactly one boundary component, because the addition or removal of boundary does not affect the cycles of G. We denote the one boundary component and its corresponding dual boundary vertex as f_∞, because it is analogous to the infinite face of a planar graph. Our algorithm for minimum cycle basis is only for G embedded on an orientable surface Σ. We conclude $\beta = 2g$.

Our algorithm is based on one of Kavitha, Mehlhorn, Michail and Paluch [27] which is in turn based on an algorithm of de Pina [12]. Our algorithm incrementally adds simple cycles $\gamma_1, \ldots, \gamma_{m-n+1}$ to the minimum cycle basis. In order to do so, it maintains a set of $(m-n+1)$-bit ***support vectors*** $S_1, \ldots, S_{m-n+1}$ with the following properties:

- The support vectors form a basis for $\mathbb{Z}_2^{m-n+1}$.
- When the algorithm is about to compute the jth simple cycle γ_j for the minimum cycle basis, $\langle S_j, [\gamma_{j'}]\rangle = 0$ for all $j' < j$.

Our algorithm chooses for γ_j the minimum weight cycle γ such that $\langle S_j, [\gamma]\rangle = 1$. Note that S_j must have at least one bit set to 1, because the set of vectors $S_1, \ldots, S_{m-n+1}$ forms a basis. Therefore, such a γ does exist; in particular, we could choose $[\gamma]$ to contain exactly one bit equal to 1 which matches any 1-bit of S_j. The correctness of choosing γ_j as above is guaranteed by the following lemma.

Lemma 4.1. *Let S be an $(m-n+1)$-bit vector with at least one bit set to 1, and let η be the minimum weight cycle such that $\langle S, [\eta]\rangle = 1$. Then, η is a member of the minimum cycle basis.*

Proof: Let $\eta_1, \ldots, \eta_{2^{m-n+1}}$ be the collection of cycles ordered by increasing weight, and choose j such that $\eta_j = \eta$. For any subset Υ of $\{\eta_1, \ldots, \eta_{j-1}\}$, we have $\langle [\bigoplus_{\eta'\in\Upsilon}\eta'], S\rangle = 0$. Therefore, η is independent of $\{\eta_1, \ldots, \eta_{j-1}\}$. Sets of independent cycles form a matroid, so η must be a member of the minimum weight cycle basis. □

4.1 Maintaining support vectors

Our algorithm updates the support vectors and computes minimum cycle basis vectors in a recursive manner. Initially, each support vector S_i has only its ith bit set to 1. Borrowing nomenclature from Kavitha *et al.* [27], we define two procedures, extend(j,k) which extends the current set of basis cycles by adding k cycles starting with γ_j, and update(j,k) which updates support vectors $S_{j+\lfloor k/2\rfloor},\ldots,S_{j+k-1}$ so that for any j',j'' with $j+\lfloor k/2\rfloor \le j' < j+k$ and $1 \le j'' < j+\lfloor k/2\rfloor$, we have $\langle S_{j'}, [\gamma_{j''}]\rangle = 0$. Our algorithm runs extend$(1, m-n+1)$ to compute the minimum cycle basis.

We implement extend(j,k) in the following manner: If $k > 1$, then our algorithm recursively calls extend$(j,\lfloor k/2\rfloor)$ to add $\lfloor k/2\rfloor$ cycles to the partial minimum cycle basis. It then calls update(j,k) so that support vectors $S_{j+\lfloor k/2\rfloor},\ldots,S_{j+k-1}$ become orthogonal to the newly added cycles of the partial basis. Finally, it computes the remaining $\lceil k/2\rceil$ basis cycles by calling extend$(j+\lfloor k/2\rfloor, \lceil k/2\rceil)$. If $k=1$, then $\langle S_j, [\gamma_{j'}]\rangle = 0$ for all $j' < j$. Our algorithm is ready to find basis cycle γ_j. We describe an $O(2^{2g}n)$ time procedure to find γ_j in Section 5.

We now describe update(j,k) in more detail. Our algorithm updates each support vector $S_{j'}$ where $j+\lfloor k/2\rfloor \le j' < j+k$. The vector $S_{j'}$ becomes $S'_{j'} = S_{j'} + \alpha_{j'0}S_j + \alpha_{j'1}S_{j+1} + \cdots + \alpha_{j'(\lfloor k/2\rfloor-1)}S_{j+\lfloor k/2\rfloor-1}$ for some set of scalar bits $\alpha_{j'0}\ldots\alpha_{j'(\lfloor k/2\rfloor-1)}$. After updating, the set of vectors $S_1,\ldots,S_{m-n+1}$ remains a basis for $\mathbb{Z}_2^{m-n+1}$ regardless of the choices for the α bits. Further, $\langle S'_{j'}, [\gamma_{j''}]\rangle = 0$ for all $j'' < j$ for all choices of the α bits, because extend(j,k) is only called after its support vectors are updated to be orthogonal to all minimum basis cycles $\gamma_1,\ldots,\gamma_{j-1}$.

However, it is non-trivial to guarantee $\langle S'_{j'}, [\gamma_{j''}]\rangle = 0$ for all j'' where $j \le j'' < j+\lfloor k/2\rfloor$. Let w^T denote the transpose of a vector w. Let

$$X = \begin{pmatrix} S_j \\ \cdots \\ S_{j+\lfloor k/2\rfloor-1} \end{pmatrix} \cdot \left([\gamma_j]^T \cdots [\gamma_{j+\lfloor k/2\rfloor-1}]^T\right)$$

and

$$Y = \begin{pmatrix} S_{j+\lfloor k/2\rfloor} \\ \cdots \\ S_{j+k-1} \end{pmatrix} \cdot \left([\gamma_j]^T \cdots [\gamma_{j+\lfloor k/2\rfloor-1}]^T\right).$$

Let $A = YX^{-1}$. Row $j'-j-\lfloor k/2\rfloor+1$ of matrix A contains exactly the bits $\alpha_{j'0}\ldots\alpha_{j'(\lfloor k/2\rfloor-1)}$ we are seeking [27, Section 4]. Matrices X, Y, and A can be computed in $O(nk^{\omega-1})$ time using fast matrix multiplication and inversion, implying the new support vectors $S'_{j+\lfloor k/2\rfloor},\ldots,S'_{j+k-1}$ can be computed in the same amount of time.

We can bound the running time of extend(j,k) using the following recurrence:

$$T(k) = \begin{cases} 2T(k/2) + O(nk^{\omega-1}) & \text{if } k > 1 \\ O(2^{2g}n) & \text{if } k = 1 \end{cases}$$

The total time spent in calls to extend(j,k) where $k>1$ is $O(nk^{\omega-1})$. The total time spent in calls to extend$(j,1)$ is $O(2^{2g}nk)$. Therefore, $T(k) = O(nk^{\omega-1} + 2^{2g}nk)$. The running time of our minimum cycle basis algorithm (after sparsifying G) is $T(O(n)) = O(n^\omega + 2^{2g}n^2)$.

5 Selecting Cycles

A *Horton cycle* is a simple cycle given by a shortest x,u-path, a shortest x,v-path, and the edge uv; in particular, the set of all Horton cycles is given by the set of $m-n+1$ elementary cycles for each

of the n shortest path trees [25]. Thus, in sparse graphs, there are $O(n^2)$ Horton cycles. A simple cycle γ of a graph G is *isometric* if for every pair of vertices $x, y \in \gamma$, γ contains a shortest x, y-path. Hartvigsen and Mardon prove that the cycles of any minimum cycle basis are all isometric [23]. Therefore, it suffices for us to focus on the set of isometric cycles to find the cycle γ_j as needed for Section 4.1. Amaldi *et al.* [2] show how to extract the set of distinct isometric cycles from a set of Horton cycles in $O(nm)$ time. Each isometric cycle is identified by a shortest path tree's root and a non-tree edge.

Here, we show that there are at most $O(2^{2g}n)$ isometric cycles in our graph of genus g (Section 5.1), and they can be partitioned into sets according to their homology classes. We can represent the isometric cycles in a given homology class using a tree that can be built in $O(n^2)$ time (Section 5.2). We then show that we can use these trees to find the minimum-cost cycle γ_j as needed for Section 4.1 in linear time per homology class of isometric cycles. We close with a discussion on how to improve the running time for computing and representing isometric cycles (Section 5.4). While these improvements do not improve the overall running time of our algorithm (by maintaining separate representations of the cycles according to their homology class, we require linear time per representation to process the support vector with respect to which γ_j is non-orthogonal; we also require $O(n^\omega)$ time to update the support vectors), it does further emphasize the bottleneck our algorithm faces in updating and representing the support vectors.

5.1 Isometric cycles in orientable surface embedded graphs

Here we prove some additional structural properties that isometric cycles have in orientable surface embedded graphs. To this end, we herein assume that shortest paths are unique. Hartvigsen and Mardon show how to achieve this assumption algorithmically when, in particular, all pairs of shortest paths are computed, as we do [23]. We first prove a generalization of the following lemma for the planar case by Borradaile, Sankowski and Wulff-Nilsen.

Lemma 5.1 (Borradaile *et al.* [6, Lemma 1.4]). *Let G be a graph in which shortest paths are unique. The intersection between an isometric cycle and a shortest path in G is a (possibly empty) shortest path. The intersection between two distinct isometric cycles γ and γ' in G is a (possibly empty) shortest path; in particular, if G is a planar embedded graph, γ and γ' do not cross.*

Lemma 5.2. *Two isometric cycles in a given homology class in a graph with unique shortest paths do not cross.*

Proof: Let γ and γ' be two isometric cycles in a given homology class. Suppose for a contradiction that γ and γ' cross. By the first part of Lemma 5.1, and the assumption that γ and γ' cross, $\gamma \cap \gamma'$ is a single simple path p. Therefore, γ and γ' cross exactly once.

Suppose γ and γ' are not null-homologous. Cutting the surface open along γ results in a connected surface with two boundary components which are connected by γ'. Cutting the surface further along γ' does not disconnect the surface. Therefore $\gamma \oplus \gamma'$ does not disconnect the surface, and so γ and γ' are not homologous, a contradiction.

If γ and γ' are null-homologous, then cutting the surface open along γ results in a disconnected surface in which $\gamma' \backslash p$ is a path, but between different components of the surface, a contradiction. □

Corollary 5.3. *There are at most ℓ distinct isometric cycles in a given homology class in a graph with ℓ faces and unique shortest paths.*

Proof: Consider the set $\{C_1, C_2, \ldots\}$ of distinct isometric cycles in a given homology class other than the null homology class. We prove by induction that $\{C_1, C_2, \ldots, C_i\}$ cut the surface into non-trivial components, each of which is bounded by exactly two of $C_1, C_2, \ldots, C_i$; this is true for C_1, C_2 since they are homologous, distinct and do cross. C_{i+1} must be contained in one component, bounded by, say, C_j and C_k since C_{i+1} does not cross any other cycle. Cutting this component along C_{i+1} creates two components bounded by C_j, C_{i+1} and C_k, C_{i+1}, respectively. Since the cycles are distinct, these component must each contain at least one face. A similar argument holds for the set of null-homologous isometric cycles. □

Since there are 2^{2g} homology classes and $\ell = O(n)$, we get:

Corollary 5.4. *There are $O(2^{2g}n)$ distinct isometric cycles in a graph of orientable genus g with unique shortest paths.*

We remark that Lemma 5.2 is not true for graphs embedded in non-orientable surfaces, because homologous cycles may cross exactly once. In fact, one can easily construct an arbitrarily large collection of homologous cycles that are pairwise crossing in a graph embedded in the projective plane.

5.2 Representing isometric cycles in each homology class

We begin by determining the homology classes of each of the $O(2^{2g}n)$ isometric cycles in the following manner. Let p be a simple path, and let $[p]_h$ denote the bitwise exclusive-or of the homology signatures of its edges. Let r be the root of any shortest path tree T. Recall, $\sigma(r, v)$ denotes the shortest path between r and v. It is straightforward to compute $[\sigma(r, v)]_h$ for every vertex $v \in V$ in $O(gn)$ time by iteratively computing signatures in a leafward order. Then, the homology signature of any isometric cycle $\gamma = \sigma(r, u) \cdot uv \cdot \sigma(v, r)$ can be computed in $O(g)$ time as $[\sigma(r, u)]_h \oplus [uv]_h \oplus [\sigma(r, v)]_h$. We spend $O(2^{2g}gn) = O(2^{2g}n^2)$ time total computing homology signatures and therefore homology classes. For the remainder of this section, we consider a set of isometric cycles $\mathcal{C}$ in a single homology class.

Let $\gamma, \gamma' \in \mathcal{C}$ be two isometric cycles in the same homology class. The combination $\gamma \oplus \gamma'$ forms the boundary of a subset of faces. That is, $(G \not\curvearrowleft \gamma) \not\curvearrowleft \gamma'$ contains at least two components. We represent the cycles in $\mathcal{C}$ by a tree $T_{\mathcal{C}}$ where each edge e of $T_{\mathcal{C}}$ corresponds to a cycle $\gamma(e) \in \mathcal{C}$ and each node v in $T_{\mathcal{C}}$ corresponds to a subset $F(v) \in (F \cup \{f_\infty\})$; specifically, the nodes correspond to sets of faces in the components of $G \not\curvearrowleft \mathcal{C}$. This tree generalizes the *region tree* defined by Borradaile, Sankowski and Wulff-Nilsen for planar graphs [6] to more general orientable surface embedded graphs. We also designate a single representative cycle $\gamma(\mathcal{C})$ of $\mathcal{C}$ and pre-compute its cycle signature $[\gamma(\mathcal{C})]$ for use in our basis cycle finding procedure. See Figure 2.

We describe here the construction of $T_{\mathcal{C}}$. Initially, $T_{\mathcal{C}}$ is a single vertex with one (looping) edge to itself (we will guarantee $T_{\mathcal{C}}$ is a tree later). Let γ_0 be an arbitrary cycle in $\mathcal{C}$. We compute $G' = G \not\curvearrowleft \gamma_0$. For the one vertex v of $T_{\mathcal{C}}$, we set $F(v) = F \cup \{f_\infty\}$ and for the one edge e, we set $\gamma(e) = \gamma_0$.

We maintain the invariants that every component of G' is bound by at least two cycles of $\mathcal{C}$ (initially the cycle γ_0 is used twice), each vertex of $T_{\mathcal{C}}$ is associated with all faces in one component of G' (possibly including f_∞), and each edge e in $T_{\mathcal{C}}$ is associated with the cycle in $\mathcal{C}$ bounding the faces for the two vertices incident to e. Assuming these invariants are maintained, and because cycles in $\mathcal{C}$ do not cross, each cycle in $\mathcal{C}$ lies entirely within some component of G'. For each cycle $\gamma \in \mathcal{C} \setminus \{\gamma_0\}$, we set $G' := G' \not\curvearrowleft \gamma$, subdivide the vertex associated with the faces of C's component,

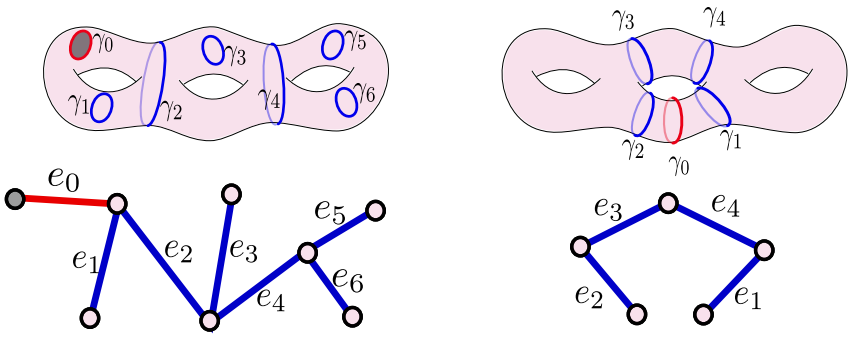

Figure 2. Two collections of homologous cycles and their generalized region trees. Left: The cycles are null-homologous. Right: The cycles lie in a non-trivial homology class.

associate the two sets of faces created in G' with the two new vertices of $T_{\mathcal{C}}$, and associate the new edge of $T_{\mathcal{C}}$ with γ.

Let r be the vertex of $T_{\mathcal{C}}$ associated with f_∞. If cycles in $\mathcal{C}$ have trivial homology, then they each separate G, and $T_{\mathcal{C}}$ is a tree. We root $T_{\mathcal{C}}$ at r and let $\gamma(\mathcal{C})$ be an arbitrary cycle. Otherwise, let e be an arbitrary edge incident to r. We set $\gamma(\mathcal{C})$ to be $\gamma(e)$, remove e from $T_{\mathcal{C}}$, and root T at r. Observe that T has exactly one leaf other than r in this case.

Computing $G' \not\leftarrow \gamma$ for one cycle γ takes $O(n)$ time. Therefore, we can compute $T_{\mathcal{C}}$ in $O(n^2)$ time total.

5.3 Selecting an isometric cycle from a homology class

Let S be an $(m-n+1)$-bit support vector. We describe a procedure to compute $\langle S, [\gamma]\rangle$ for every isometric cycle γ in G in $O(2^{2g}n)$ time. Using this procedure, we can easily return the minimum weight cycle such that $\langle S, [\gamma]\rangle = 1$.

We begin describing the procedure for cycles in the trivial homology class. Let $\mathcal{C}$ be the collection of null-homologous isometric cycles computed above, and let $T_{\mathcal{C}}$ be the tree computed for this set. Consider any edge e of $T_{\mathcal{C}}$. The first $2g$ bits of $[\gamma(e)]$ are equal to 0, because any co-cycle crosses a cut in the dual an even number of times. Cycle $\gamma(e)$ bounds a subset of faces F'. In particular, F' is the set of faces associated with vertices lying *below* e in $T_{\mathcal{C}}$. The ith bit of $[\gamma(e)]$ is 1 if and only if p_i crosses cut F' an odd number of times; in other words, the ith bit is 1 if and only if $f_i \in F'$.

We compute $\langle S, [\gamma]\rangle$ for every cycle $\gamma \in \mathcal{C}$ in $O(n)$ time by essentially walking up $T_{\mathcal{C}}$ in the following manner. For each edge e in $T_{\mathcal{C}}$ going to a leaf v, we maintain a bit z initially equal to 0 and iterate over each face $f_i \in F(v)$. If the ith bit of S is equal to 1 then we flip z. After going through all the faces in $F(v)$, z is equal to $\langle S, [\gamma(e)]\rangle$.

We then iterate up the edges of $T_{\mathcal{C}}$ toward the root. For each edge e, we let v be the lower endpoint of e and set bit z equal to the exclusive-or over all $\langle S, \gamma(e')\rangle$ for edges e' lying below v. We then iterate over the faces of $F(v)$ as before and set $\langle S, [\gamma(e)]\rangle$ equal to z as before. We iterate over every face of G at most once during this procedure, so it takes $O(n)$ time total.

Now, consider the set of isometric cycles $\mathcal{C}$ for some non-trivial homology class. Consider any edge e of $T_{\mathcal{C}}$. Let F' be the subset of faces bound by $\gamma(\mathcal{C}) \oplus \gamma(e)$. The ith bit of $[\gamma(e)]$ disagrees with the ith bit of $[\gamma(\mathcal{C})]$ if and only if path p_i crosses dual cut F' an odd number of times; in other words, the ith bits differ if and only if $f_i \in F'$. By construction, $\gamma(\mathcal{C})$ lies on the boundary of $F(r)$ and $F(v)$ where r and v are the root and other leaf of $T_{\mathcal{C}}$ respectively. Root r is the only node of $T_{\mathcal{C}}$ associated with f_∞. We conclude the ith bit of $[\gamma(e)]$ disagrees with $[\gamma(\mathcal{C})]$ if and only if f_i is associated with a vertex lying below e in $T_{\mathcal{C}}$.

We again walk up $T_{\mathcal{C}}$ to compute $\langle S, [\gamma] \rangle$ for every cycle $\gamma \in \mathcal{C}$. Recall, $[\gamma(\mathcal{C})]$ is precomputed and stored with $T_{\mathcal{C}}$. For each edge e of $T_{\mathcal{C}}$ in rootward order, let v be the lower endpoint of e. Let e' be the edge lying below e in $T_{\mathcal{C}}$ if it exists. If e' does not exist, we denote $\gamma(e')$ as $\gamma(\mathcal{C})$. We set z equal to $\langle S, \gamma(e') \rangle$. We then iterate over the faces of $F(v)$ as before, flipping z once for every bit i where $f_i \in F(v)$ and bit i of S is equal to 1. We set $\langle S, \gamma(e) \rangle := z$. As before, we consider every face at most once, so the walk up $T_{\mathcal{C}}$ takes $O(n)$ time.

We have shown the following lemma, which concludes the discussion of our minimum cycle basis algorithm.

Lemma 5.5. *Let G be a graph with n vertices, m edges, and ℓ faces cellulary embedded in an orientable surface of genus g such that $m = O(n)$ and $\ell = O(n)$. We can preprocess G in $O(2^{2g}n^2)$ time so that for any $(m-n+1)$-bit support vector S we can compute the minimum weight cycle γ such that $\langle S, \gamma \rangle = 1$ in $O(2^{2g}n)$ time.*

Theorem 5.6. *Let G be a graph with n vertices and m edges, cellularly embedded in an orientable surface of genus g. We can compute a minimum weight cycle basis of G in $O(n^\omega + 2^{2g}n^2 + m)$ time.*

5.4 Improving the time for computing and representing isometric cycles

Here we discuss ways in which we can improve the running time for finding and representing isometric cycles using known techniques, thereby isolating the bottleneck of the algorithm to updating the support vectors and computing γ_j.

The set and representation of isometric cycles can computed recursively using $O(\sqrt{gn})$ balanced separators (e.g. [1]) as inspired by Wulff-Nilsen [36]. Briefly, given a set S of $O(\sqrt{gn})$ separator vertices (for a graph of bounded genus), find all the isometric cycles in each component of $G \setminus S$ and represent these isometric cycles in at most 2^{2g} region trees per component, as described above. Merging the region trees for different components of $G \setminus S$ is relatively simple since different sets of faces are involved. It remains to compute the set of isometric cycles that contain vertices of S and add them to their respective region trees. First note that a cycle that is isometric in G and does not contain a vertex of S is isometric in $G \setminus S$, but a cycle that is isometric in $G \setminus S$ may not be isometric in G, so indeed we are computing a superset of the set of isometric cycles via this recursive procedure. However, it is relatively easy to show that an isometric cycle of $G \setminus S$ can cross an isometric cycle of G at most once, so, within a given homology class, isometric cycles will nest and be representable by a region tree.

To compute the set of isometric cycles that intersect vertices of S, we first compute shortest path trees rooted at each of the vertices of S, generating the Horton cycles rooted at these vertices; this procedure takes $O(\sqrt{gn} \cdot n)$ time using the linear time shortest path algorithm for graphs excluding minors of sub-linear size [34]. We point out that the algorithm of Amaldi *et al.* [2] works by identifying Horton cycles that are not isometric and by identifying, among different Horton-cycle representations of a given isometric cycle, one representative; this can be done for a subset of Horton cycles, such as those rooted in vertices of S, and takes time proportional to the size of the representation of the Horton cycles (i.e., the $O(\sqrt{n})$ shortest path trees, or $O(\sqrt{g}n^{1.5})$).

For a given homology class of cycles, using the shortest-path tree representation of the isometric cycles, we can identify those isometric cycles in that homology class by computing the homology signature of root-to-node paths in the shortest path tree as before; this process can be done in $O(\sqrt{g}n^{1.5})$ time. We must now add these cycles to the corresponding region tree. Borradaile, Sankowski and Wulff-Nilsen [6] describe a method for adding n cycles to a region tree in $O(n \operatorname{poly} \log n)$ time that is used in their minimum cycle basis algorithm for planar graphs; this method will generalize to surfaces for nesting cycles. Therefore computing the homology classes

of these isometric cycles and adding these isometric cycles to the region trees takes a total of $O(2^{2g}\sqrt{g}n^{1.5})$ time.

In total, this recursive method for computing and building a representation of a superset of the isometric cycles takes time given by the recurrence relation $T(n) = 2T(n/2) + O(2^{2g}\sqrt{g}n^{1.5})$ or $O(2^{2g}\sqrt{g}n^{1.5})$ time.

6 Homology Basis

We now describe our algorithm for computing a minimum homology basis. Our algorithm works for both orientable and non-orientable surfaces, although we assume without loss of generality that the surface contains at least one boundary component. At a high level, our algorithm for minimum homology bases is very similar to our algorithm for minimum cycle bases. As before, our algorithm incrementally adds simple cycles $\gamma_1, \ldots, \gamma_\beta$ to the minimum homology basis by maintaining a set of β support vectors $S_1, \ldots, S_\beta$ such that the following hold:

- The support vectors form a basis for $\mathbb{Z}_2^\beta$.
- When the algorithm is about to compute the jth cycle γ_j for the minimum homology basis, $\langle S_j, [\gamma_{j'}]_h \rangle = 0$ for all $j' < j$.

Our algorithm chooses for γ_j the minimum weight simple cycle γ such that $\langle S_j, [\gamma]_h \rangle = 1$. The following lemma has essentially the same proof as Lemma 4.1.

Lemma 6.1. *Let S be a β-bit vector with at least one bit set to 1, and let η be the minimum weight cycle such that $\langle S, [\eta]_h \rangle = 1$. Then, η is a member of the minimum homology basis.*

As before, our algorithm updates the support vectors and computes minimum homology basis cycles in a recursive manner. We define extend(j, k) and update(j, k) as before, using homology signatures in place of cycle signatures when applicable. Our algorithm runs extend$(1, \beta)$ to compute the minimum homology basis.

The one crucial difference between our minimum cycle basis and minimum homology basis algorithms is the procedure we use to find each minimum homology basis cycle γ_j given support vector S_j. The homology basis procedure takes $O(\beta^2 n \log n)$ time instead of $O(2^{2g}n)$ time, and it requires no preprocessing step. We describe the procedure in Sections 6.1 and 6.2.

The procedure update(j, k) takes only $O(\beta k^{\omega-1})$ time in our minimum homology basis algorithm, because signatures have length β. Therefore, we can bound the running time of extend(j, k) using the following recurrence:

$$T(k) = \begin{cases} 2T(k/2) + O(\beta k^{\omega-1}) & \text{if } k > 1 \\ O(\beta^2 n \log n) & \text{if } k = 1 \end{cases}$$

The total time spent in calls to extend(j, k) where $k > 1$ is $O(\beta k^{\omega-1})$. The total time spent in calls to extend$(j, 1)$ is $O(\beta^2 kn \log n)$. Therefore, $T(k) = O(\beta k^{\omega-1} + \beta^2 kn \log n)$. The running time of our minimum homology basis algorithm[3] (after sparsifying G) is $T(\beta) = O(\beta^3 n \log n) = O((g+b)^3 n \log n)$.

[3] Our minimum homology basis algorithm can be simplified somewhat by having extend(j, k) recurse on extend$(j, 1)$ and extend$(j+1, k-1)$ and by using a simpler algorithm for update(j, k). This change will increase the time spent in calls to extend(j, k) where $k > 1$, but the time taken by calls with $k = 1$ will still be a bottleneck on the overall run time.

6.1 Cyclic double cover

In order to compute minimum homology basis cycle γ_j, we lift the graph into a *covering space* known as the ***cyclic double cover***. Our presentation of the cyclic double cover is similar to that of Erickson [16]. Erickson describes the cyclic double cover relative to a single simple non-separating cycle in an orientable surface; however, we describe it relative to an arbitrary *set* of non-separating *co-paths* determined by a support vector S, similar to the homology cover construction of Erickson and Nayyeri [17]. Our construction works for non-orientable surfaces without any special considerations.

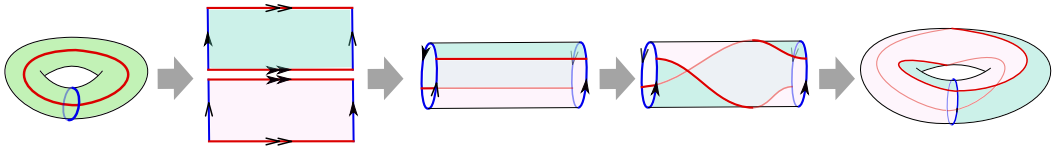

Figure 3. Constructing the cyclic double cover. Left to right: A pair of co-cycles Ψ on the torus Σ; the surfaces $(\Sigma', 0)$ and $(\Sigma', 1)$; identifying copies of one co-cycle; preparing to identify copies of the other co-cycle; the cyclic double cover.

Let S be a β-bit support vector for the minimum homology basis problem as defined above. We define the cyclic double cover relative to S using a standard *voltage construction* [22, Chapter 4]. Let G_S^2 be the graph whose vertices are pairs (v, z), where v is a vertex of G and z is a bit. The edges of G_S^2 are ordered pairs $(uv, z) := (u, z)(v, z \oplus \langle S, [uv]_h \rangle)$ for all edges uv of G and bits z. Let $\pi : G_S^2 \to G$ denote the ***covering map*** $\pi(v, z) = v$. The ***projection*** of any vertex, edge, or path in G_S^2 is the natural map to G induced by π. We say a vertex, edge, or path p in G ***lifts*** to p' if p if the projection of p'. A closed path in G_S^2 is defined to bound a face (be a boundary component) of G_S^2 if and only if its projection with regard to π bounds a face (is a boundary component) of G. This construction defines an embedding of G_S^2 onto a surface Σ_S^2 (we will prove G_S^2 and Σ_S^2 are connected shortly).

We can also define G_S^2 in a more topologically intuitive way as follows. Let Ψ be a set of co-paths which contains each co-path p_i for which the ith bit of S is equal to 1. Let Σ' be the surface obtained by cutting Σ along the image of each co-path in Ψ. Note that Σ' may be disconnected. Each co-path $p_i \in \Psi$ appears as two copies on the boundary of Σ' denoted p_i^- and p_i^+ (note that p_i^- and p_i^+ may themselves be broken into multiple components). Create two copies of Σ' denoted $(\Sigma', 0)$ and $(\Sigma', 1)$, and let (p_i^-, z) and (p_i^+, z) denote the copies of p_i^- and p_i^+ in surface (Σ', z). For each co-path $p_i \in \Psi$, we identify $(p_i^+, 0)$ with $(p_i^-, 1)$ and we identify $(p_i^+, 1)$ with $(p_i^-, 0)$, creating the surface Σ_S^2 and the graph G_S^2 embedded on Σ_S^2. See Figure 3.

The first three of the following lemmas are immediate.

Lemma 6.2. *Let γ be any simple cycle in G, and let s be any vertex of γ. Then γ is the projection of a unique path in G_S^2 from $(s, 0)$ to $(s, \langle S, [\gamma]_h \rangle)$.*

Lemma 6.3. *Every lift of a shortest path in G is a shortest path in G_S^2.*

Lemma 6.4. *Let γ be the minimum weight simple cycle of G such that $\langle S, [\gamma]_h \rangle = 1$, and let s be any vertex of γ. Then γ is the projection of the shortest path in G_S^2 from $(s, 0)$ to $(s, 1)$.*

Lemma 6.5. *The cyclic double cover G_S^2 is connected.*

Proof: There exists some simple cycle γ in G such that $\langle S, [\gamma]_h \rangle = 1$. Let s be any vertex of γ. Let v be any vertex of G. We show there exists a path from (v, z) to $(s, 0)$ in G_S^2 for both bits z.

There exists a path from v to s in G so there is a path from (v, z) to one of $(s, 0)$ or $(s, 1)$ in G_S^2. The other of $(s, 0)$ or $(s, 1)$ may be reached by following the lift of γ. □

Observe that G_S^2 has $2n$ vertices and $2m$ edges. Each co-path p_i shares an even number of edges with each face of G. By Lemma 6.2, both lifts of any face f to G_S^2 are cycles; in particular both lifts are faces. However, there may be one or more boundary cycles γ of G such that $\langle S, [\gamma]_h \rangle = 1$. It takes both lifts of such a cycle γ to make a single boundary component in G_S^2. We conclude G_S^2 contains 2ℓ faces and between b and $2b$ boundary cycles. Surface Σ_S^2 has Euler characteristic $2n - 2m + 2\ell = 2\chi$. It is non-orientable if and only if there exists a one-sided cycle η such that $\langle S, [\eta]_h \rangle = 0$. If both Σ and Σ_S^2 are non-orientable, then Σ_S^2 has genus at most $2g + b$. If only Σ is non-orientable, then Σ_S^2 has genus at most $g + b/2 - 1$. If both surfaces are orientable, then Σ_S^2 has genus at most $2g + b/2 - 1$. In all three cases, the genus is at most $O(\beta)$.

6.2 Selecting homology basis cycles

Let S be any β-bit support vector. We now describe our algorithm to select the minimum weight cycle γ such that $\langle S, [\gamma]_h \rangle = 1$. Our algorithm is based on one by Erickson and Nayyeri [17] for computing minimum weight cycles in arbitrary homology classes, except we use the cyclic double cover instead of their $\mathbb{Z}_2$-homology cover. We have the following lemma. While it was shown with orientable surfaces in mind, the proof translates verbatim to the non-orientable case.

Lemma 6.6 (Erickson and Nayyeri [17, Lemma 5.1]). *In $O(n \log n + \beta n)$ time, we can construct[4] a set Π of $O(\beta)$ shortest paths in G, such that every non-null-homologous cycle in G intersects at least one path in Π.*

Let G_S^2 be the cyclic double cover of G with regard to S. Our algorithm constructs G_S^2 in $O(\beta n)$ time.

Suppose our desired cycle γ intersects shortest path $\sigma \in \Pi$ at some vertex s. By Lemma 6.4, simple cycle γ is the projection of the shortest path in G_S^2 from $(s, 0)$ to $(s, 1)$. Let $\hat{\gamma}$ be this shortest path in G_S^2. Let $\hat{\sigma}$ be the lift of σ to G_S^2 that contains vertex $(s, 0)$. By Lemma 6.3, path $\hat{\sigma}$ is also a shortest path in G_S^2. If $\hat{\gamma}$ uses any other vertex (v, z) of $\hat{\sigma}$ other than $(s, 0)$, then it can use the entire subpath of $\hat{\sigma}$ between $(s, 0)$ and (v, z).

Now, consider the surface $\Sigma_S^2 \not\curvearrowright \hat{\sigma}$ which contains a single face bounded by two copies of $\hat{\sigma}$ we denote $\hat{\sigma}^-$ and $\hat{\sigma}^+$. For each vertex (v, z) on $\hat{\sigma}$, let $(v, z)^-$ and $(v, z)^+$ denote its two copies on $\hat{\sigma}^-$ and $\hat{\sigma}^+$ respectively. From the above discussion, we see $\hat{\gamma}$ is a shortest path in $\Sigma_S^2 \not\curvearrowright \hat{\sigma}$ from one of $(s, 0)^-$ or $(s, 0)^+$ to $(s, 1)$.

To find γ, we use the following generalization of Klein's [29] multiple-source shortest path algorithm:

Lemma 6.7 (Cabello *et al.* [8]). *Let G be a graph with n vertices, cellularly embedded in a surface of genus g, and let f be any face of G. We can preprocess G in $O(gn \log n)$ time and $O(n)$ space so that the length of the shortest path from any vertex incident to f to any other vertex can be retrieved in $O(\log n)$ time.*

Our algorithm iterates over the $O(\beta)$ shortest paths present in Π. For each such path σ, it computes a lift $\hat{\sigma}$ in G_S^2, cuts Σ_G^2 along $\hat{\sigma}$, and runs the multiple-source shortest path procedure of Lemma 6.7 to find the shortest path from some vertex $(s, z)^{\pm}$ on $\hat{\sigma}^{\pm}$ to $(s, z \oplus 1)$. Each shortest path

[4] We only need to construct Π once for the entire minimum homology basis algorithm, but constructing it once per basis cycle does not affect the overall run time.

it finds projects to a closed path γ' such that $\langle S, [\gamma']_h \rangle = 1$. By the above discussion, the shortest such projection can be chosen for γ. Running the multiple-source shortest path procedure $O(\beta)$ times on a graph of genus $O(\beta)$ takes $O(\beta^2 n \log n)$ time total. We conclude the discussion of our minimum weight homology basis algorithm.

Lemma 6.8. *Let G be a graph with n vertices, m edges, and ℓ faces cellulary embedded in a surface of genus g such that $m = O(n)$ and $\ell = O(n)$. For any β-bit support vector S we can compute the minimum weight cycle γ such that $\langle S, \gamma \rangle = 1$ in $O(\beta^2 n \log n)$ time.*

Theorem 6.9. *Let G be a graph with n vertices and m edges, cellularly embedded in an orientable or non-orientable surface of genus g with b boundary components. We can compute a minimum weight homology basis of G in $O((g+b)^3 n \log n + m)$ time.*

References

[1] L. Aleksandrov and H. Djidjev. Linear algorithms for partitioning embedded graphs of bounded genus. *SIAM J. of Disc. Math.* 9(1):129–150, 1996.

[2] E. Amaldi, C. Iuliano, T. Jurkiewicz, K. Mehlhorn, and R. Rizzi. Breaking the $O(m^2 n)$ barrier for minimum cycle bases. *Proc. 17th Ann. Euro. Symp. Algo.*, pp. 301–312. 2009.

[3] F. Berger, P. Gritzmann, and S. de Vries. Minimum cycle bases for network graphs. *Algorithmica* 40(1):51–62, 2004.

[4] G. Borradaile, E. W. Chambers, K. Fox, and A. Nayyeri. Minimum cycle and homology bases of surface embedded graphs. *Proc. 32nd Ann. Int. Symp. Comput. Geom.*, pp. 23:1–23:15. 2016.

[5] G. Borradaile, D. Eppstein, A. Nayyeri, and C. Wulff-Nilsen. All-pairs minimum cuts in near-linear time for surface-embedded graphs. *Proc. 32nd Ann. Int. Symp. Comput. Geom.*, pp. 22:1–22:16. 2016.

[6] G. Borradaile, P. Sankowski, and C. Wulff-Nilsen. Min st-cut oracle for planar graphs with near-linear preprocessing time. *ACM Trans. Algo.* 11(3):16, 2015.

[7] O. Busaryev, S. Cabello, C. Chen, T. K. Dey, and Y. Wang. Annotating simplicies with a homology basis and its applications. *Proc. 13th Scandinavian Workshop on Algo. Theory*, pp. 189–200. 2012.

[8] S. Cabello, E. W. Chambers, and J. Erickson. Multiple-source shortest paths in embedded graphs. *SIAM J. Comput.* 42(4):1542–1571, 2013.

[9] A. C. Cassell, J. C. de Henderson, and K. Ramachandran. Cycle bases of minimal measure for the structural analysis of skeletal structures by the flexibility method. *Proc. R. Soc. Lond. Ser. A* 350:61–70, 1976.

[10] L. O. Chua and L.-K. Chen. On optimally sparse cycle and coboundary basis for a linear graph. *IEEE Trans. Circuit Theory* 20:495–503, 1973.

[11] É. Colin de Verdière. Topological algorithms for graphs on surfaces. Habilitation thesis, May 2012.

[12] J. C. de Pina. *Applications of Shortest Path Methods.* Ph.D. thesis, University of Amsterdam, 1995.

[13] T. K. Dey, J. Sun, and Y. Wang. Approximating loops in a shortest homology basis from point data. *Proc. 26th Ann. Symp. Comput. Geom.*, pp. 166–175. 2010.

[14] H. Edelsbrunner and J. Harer. *Computational Topology, An Introduction.* American Mathematical Society, 2010.

[15] D. Eppstein. Dynamic generators of topologically embedded graphs. *Proc. 14th Ann. ACM-SIAM Symp. Disc. Algo.*, pp. 599–608. 2003.

[16] J. Erickson. Shortest non-trivial cycles in directed surface graphs. *Proc. 27th Ann. Symp. Comput. Geom.*, pp. 236–243. 2011.

[17] J. Erickson and A. Nayyeri. Minimum cuts and shortest non-separating cycles via homology covers. *Proc. 22nd Ann. ACM-SIAM Symp. Disc. Algo.*, pp. 1166–1176. 2011.

[18] J. Erickson and K. Whittlesey. Greedy optimal homotopy and homology generators. *Proc. 16th Ann. ACM-SIAM Symp. on Disc. Algo.*, pp. 1038–1046. 2005.

[19] K. Fox. Shortest non-trivial cycles in directed and undirected surface graphs. *Proc. 24th Ann. ACM-SIAM Symp. Disc. Algo.*, pp. 352–364. 2013.

[20] A. Golynski and J. D. Horton. A polynomial time algorithm to find the minimum cycle basis of a regular matroid. *Proc. 8th Scandinavian Workshop on Algo. Theory*, pp. 200–209. 2002.

[21] R. E. Gomory and T. C. Hu. Multi-terminal network flows. *J. SIAM* 9(4):551–570, 1961.

[22] J. L. Gross and T. W. Tucker. *Topological graph theory.* Dover Publications, 2001.

[23] D. Hartvigsen and R. Mardon. The all-pairs min cut problem and the minimum cycle basis problem on planar graphs. *SIAM J. Disc. Math.* 7(3):403–418, 1994.

[24] A. Hatcher. *Algebraic Topology.* Cambridge Univ. Press, 2002. ⟨http://www.math.cornell.edu/~hatcher/AT/ATpage.html⟩.

[25] J. D. Horton. A polynomial-time algorithm to find the shortest cycle basis of a graph. *SIAM J. Comput.* 16(2):358–366, 1987.

[26] G. F. Italiano, Y. Nussbaum, P. Sankowski, and C. Wulff-Nilsen. Improved algorithms for min cut and max flow in undirected planar graphs. *Proc. 43rd Ann. ACM Symp. Theory Comput.*, pp. 313–322. 2011.

[27] T. Kavitha, K. Mehlhorn, D. Michail, and K. E. Paluch. An $\tilde{O}(m^2n)$ algorithm for minimum cycle basis of graphs. *Algorithmica* 52(3):333–349, 2008.

[28] G. Kirchhoff. Ueber die auflösung der gleichungen, auf welche man bei der untersuchung der linearen vertheilung galvanischer ströme geführt wird. *Poggendorf Ann. Physik* 72:497–508, 1847. English transl. in Trans. Inst. Radio Engrs. CT-5 (1958), pp. 4–7.

[29] P. Klein. Multiple-source shortest paths in planar graphs. *Proc. 16th Ann. ACM-SIAM Symp. Disc. Algo.*, pp. 146–155. 2005.

[30] D. E. Knuth. *The Art of Computer Programming.* vol. 1. Addison-Wesley, 1968.

[31] K. Mehlhorn and D. Michail. Minimum cycle bases: Faster and simpler. *ACM Trans. Algo.* 6(1):8, 2009.

[32] B. Mohar and C. Thomassen. *Graphs on Surfaces.* Johns Hopkins Univ. Press, 2001.

[33] J. R. Munkres. *Topology*, 2nd edition. Prentice-Hall, 2000.

[34] S. Tazari and M. Müller-Hannemann. Shortest paths in linear time on minor-closed graph classes with an application to Steiner tree approximation. *Disc. Applied Math.* 157(4):673–684, 2009.

[35] G. Tewari, C. Gotsman, and S. J. Gortler. Meshing genus-1 point clouds using discrete one-forms. *Comput. Graph.* 30(6):917–926, 2006.

[36] C. Wulff-Nilsen. Minimum cycle basis and all-pairs min cut of a planar graph in subquadratic time. Tech. Rep. arXiv:0912.1208, University of Copenhagen, 2009.

[37] A. Zomorodian. *Topology for Computing.* Cambridge Univ. Press, 2005.